New frontiers in quantum chemistry using supercomputers

Jeff Hammond

Leadership Computing Facility Argonne National Laboratory

15 March 2011



Computation and experiment

W. Graham Richards (Oxford) declared the "third age of quantum chemistry" in 1979, noting that computation caught up to experiment and that synergy is essential.

In 1979, one of the pioneers of HPC in quantum chemistry, Fritz Schaefer, shared time on the University of Texas CDC Cyber 170/750 running at 1.6 Mflops.

My iPhone 3Gs will run at \sim 6 Mflops for the SciMark 2.0 benchmark.

Unfortunately, the experimentalists move at better than a Moore's law. . .

Three pillars

The US DOE has declared that simulation was one of three pillars of science, along with theory and experiment.

At least in the US, convincing the theorists and experimentalists that computational work is as valuable as theirs is not easy.

It is much easy to justify computational science when the experiments involve setting off nuclear weapons, but we need more general arguments for the utility of simulation.

Overview of Computational Chemistry

Atomistic simulation in chemistry

- classical molecular dynamics (MD) with empirical potentials
- ab initio molecular dynamics based upon density-function theory (DFT)
- quantum chemistry with wavefunctions
 e.g. perturbation theory (PT), Coupled-Cluster
 (CC) or Quantum Monte Carlo (QMC).

Classical molecular dynamics

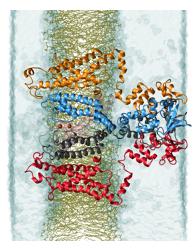


Image courtesy of Benoît Roux via ALCF.

 Solves Newton's equations of motion with empirical terms and classical electrostatics.

• Size: 100K-10M atoms

• Time: 1-10 ns/day

• Scaling: $\sim N_{atoms}$

Data from K. Schulten, et al. "Biomolecular modeling in the era of petascale computing." In D. Bader, ed., *Petascale Computing:*Algorithms and Applications.

Car-Parrinello molecular dynamics

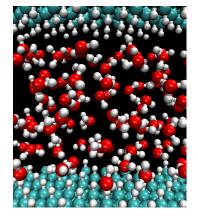


Image courtesy of Giulia Galli via ALCF.

 Forces obtained from solving an approximate single-particle
 Schrödinger equation;
 time-propagation via Lagrangian approach.

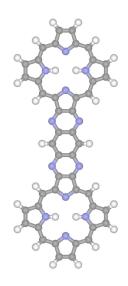
• Size: 100-1000 atoms

Time: 0.01-1 ps/day

• Scaling: $\sim N_{el}^{x}$ (x=1-3)

F. Gygi, IBM J. Res. Dev. 52, 137 (2008); E. J. Bylaska et al. J. Phys.: Conf. Ser. 180, 012028 (2009).

Wavefunction theory



- MP2 is second-order PT and is accurate via magical cancellation of error.
- CC is infinite-order solution to many-body Schrödinger equation truncated via clusters.
- QMC is Monte Carlo integration applied to the Schrödinger equation.
- Size: 10-100 atoms, maybe 100-1000 atoms with MP2.
- Time: N/A
- Scaling: $\sim N_{bf}^{x}$ (x=2-7)

Image courtesy of Karol Kowalski and Niri Govind.



Motivation for HPC

- Electronic excited-states and electric-field perturbations push the limits of conventional approximations in DFT and are outside the scope of classical methods.
- Interesting chemical processes in biology and material science require model systems too large for a conventional computational resources.
- Answering many chemical questions requires large data sets which cannot be obtained in a reasonable amount time if done sequentially.

Electronic excited-states in biology

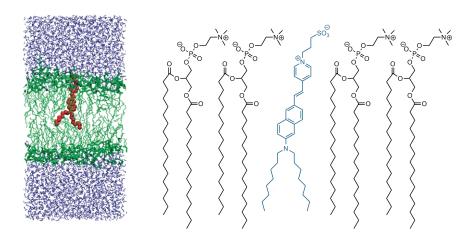
Joint work with Benoît Roux (UC/ANL) and Karol Kowalski (PNNL).

Molecular probes

Optical potentiometric probes have become important tools in electrophysiology. These organic molecules display spectroscopic responses to membrane potential and have been used for the study and characterization of model membranes, nerve and muscle tissues, organelles, microorganisms, and red blood cells. They can often be used in place of conventional microelectrodes and lend themselves to many system not accessible to microelectrodes.

E. Fluhler, V. G. Burnham, L. M. Loew, *Biochemistry* **24**, 5749 (1985). "Spectra, membrane binding, and potentiometric responses of new charge shift probes."

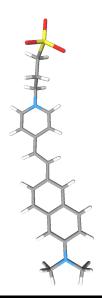
Membrane configuration of di-8-ANEPPS

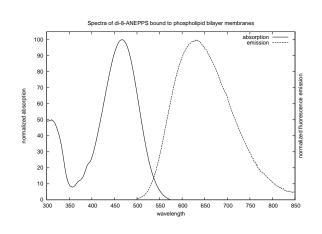


C. F. Rusu, H. Lanig, O. G. Othersen, C. Kryschi, and T. Clark, J. Phys. Chem. B 112, 2445 (2008).

ANEPPS model structure

peaks $= 2.655 \ \text{eV}$ (3.987 eV) and 1.965 eV





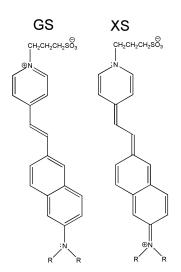
http://probes.invitrogen.com/media/spectra/data/3167lip.txt

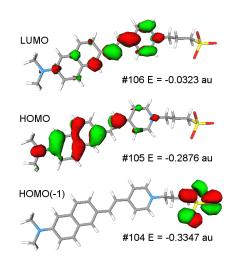
Computing the spectrum I

Root	au	eV	nm	OS
B3LYP/cc-pVDZ	0.002	0.06	19758.1	0.001
B3LYP/aug-cc-pVDZ	0.016	0.44	2799.2	0.000
PBE0/aug-cc-pVDZ	0.026	0.70	1773.5	0.000
BH&H/aug-cc-pVDZ	0.091	2.47	501.7	0.000
TDHF/aug-cc-pVDZ	0.124	3.38	366.7	1.731
CIS/aug-cc-pVDZ	0.132	3.59	345.8	1.949

The basis set dependence is an illusion.

What are the electrons doing?





Computing the spectrum II

Basis	au	eV	au	nm
CC2/6-31G*	1.5	1.815	0.067	683.3
CCSD/6-31G*	0.5	3.629	0.133	341.6
CCSD/6-31G*	0.7	3.231	0.119	383.8
CCSD/6-31G*	1.5	2.984	0.110	415.5
CCSD/6-31G*	2.0	2.962	0.109	418.6
CCSD/6-31G*	∞	2.968	0.109	417.7
CCSD/cc-pVDZ	∞	2.945	0.108	421.0
Experiment		2.655	0.098	467

For τ >1.5, X_1 is dominated by one amplitude. $||X_2||$ is small throughout.

Accurate many-body methods

Method	Basis	au	eV	au	nm
	6-31G*	0.5	3.629	0.133	341.6
EOM-CCSD	6-31G*	0.7	3.231	0.119	383.8
	6-31G*	1.5	2.980	0.110	416.1
	6-31G*	0.5	3.590	0.132	345.4
CR-EOM-CCSD(T)	6-31G*	0.7	3.150	0.116	393.6
	6-31G*	1.5	2.810	0.103	441.2
Experiment			2.655	0.098	467

Active-space approach:

Only orbitals with $-\tau < \epsilon < \tau$ are correlated. $\tau = 1.5$ corresponds to frozen-core approximation.

NWChem implementation (TCE)

Timings on 256 nodes of Chinook:

Procedure	wall time (s)
SCF total time	57
four-index transformation	192
one CCSD iteration	157
one EOM-CCSD iteration	252
CR-EOM-CCSD(T) evaluation	6301
Total time	12510

Even though this calculation is trivial with NWChem, it is still impossible with single-node codes because of the memory wall.

Karol is up to 25K cores at NERSC...

Electric properties

Method comparison

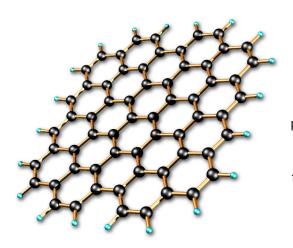
Property	E_{xs}		μ	$eta_{ }$
Unit	(nm)	(eV)	(D)	(10^{-30} esu)
HF/CIS	346	3.59	31.3	-97.1
CCSD	416	2.98	29.8	-191.8
CR(T)	441	2.81	-	-
Experiment	467			
BH&H	502	2.47	32.1	-146.4
PBE0	1773	0.70	32.6	267.9
B3LYP	2799	0.44	32.8	925.9

From what one sees for PNA, triples (e.g. CC3) will be necessary to compute $\beta(\omega_0; \omega_1, \omega_2)$.

Bottom-up simulation in material science

Joint work with Karol Kowalski (PNNL).

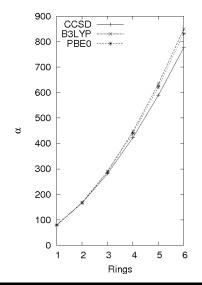
Graphitic materials



Polarizability simultaneously probes excited-state behavior (poles) and intermolecular forces — dispersion closely related to $\alpha(\omega)$.

Image from Berkeley Labs (Lanzara Group).

Polarizabilities of polyacenes

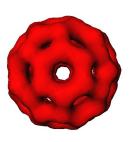


		$lpha_{ extcolor{LL}}$	
Rings	CCSD	B3LYP	PBE0
1	80.57	79.38	78.75
2	166.61	168.59	166.48
3	281.60	291.56	287.07
4	423.83	447.60	439.52
5	589.97	634.65	622.40
6	776.83	849.55	831.79

J. Chem. Phys. 127, 144105 (2007).

Polarizabilities of C₆₀

	Wavelength (nm)		
Method	∞	1064	
Lowest found	441.3	-	
B3LYP/6-31G*	469.0	-	
HF/6-31++G	506.8	515.6	
Experiment	516.3	533.1	
CCSD/Z3Pol	555.3	564.9	
$LDA/TZP{++}$	571.6	-	
CC2/6-31++G	586.8	600.8	
CC2/6-31++G*	606.8	622.6	
CC2/aug-cc-pVDZ	623.7	640.2	
Highest found	1033.2	_	



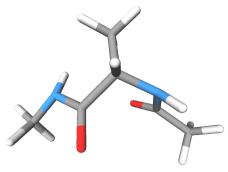
3.7 speed-up from 256 to 1024 processors.

J. Chem. Phys. 129, 226101 (2008).

Force-fields from first-principles

Joint work with Karl Freed (UC), Benoît Roux (UC/ANL), Alex MacKerell (Maryland)

The protein prototype — dialanine



- Debatable if dialanine represents the real torsional potential.
- Many FF potentials use MP2 dialanine results.
- Useful for calibrating methods without pollution of cooperative effects.
- Computationally tractable for CCSD(T) (whole ϕ - ψ map).

Evaluating models with CCSD(T)

- Used OPLS-AA geometries to prevent bias.
- 6-311++G** basis set (aug-cc-pVTZ desirable).
- Difference between $6-31+G^*$ and $6-311++G^{**}$: SCF=1.18, MP2=2.48 (MUE in kJ/mol).
- Difference between 6-311++G** and aug-cc-pVTZ: SCF=0.89, MP2=1.70 (MUE in kJ/mol).
- 350 configurations (30° grid everywhere, 10° in basins).
- CCSD(T)/6-311++G** takes approximately 1 hour per job on 64 nodes (whole set could run in an afternoon).



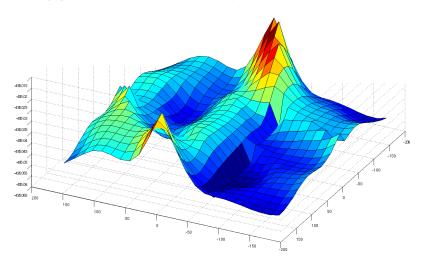
Evaluating models with CCSD(T)

Method	MUE	Max	Method	MUE	Max
MP2	0.436	1.304	M05	1.551	8.389
CCSD	0.577	1.426	TPSS+D	1.593	9.258
B2PLYP	0.913	4.690	FT97	1.711	11.047
M06	1.137	4.874	CAMB3LYP	1.747	6.268
Becke97 + D	1.177	5.981	M06-2X	1.757	5.812
Becke98	1.287	7.526	BB1K	1.773	7.310
TPSS	1.312	10.691	B2LYP	1.913	6.514
B3LYP+D	1.327	6.379	HCTH120	2.119	10.141
TPSSh	1.330	9.525	ВОР	2.614	9.118
M06-L	1.378	6.657	M06-HF	2.884	12.286
Becke97	1.391	7.486	SCF	3.066	11.076
PBE+D	1.404	9.812	HCTH407	3.168	9.678
X3LYP	1.430	7.747	HCTH	3.330	9.788
B3LYP	1.456	7.884	CAMPBE0	3.348	10.676
PBE0	1.506	8.041	(All in NWChem 6.0)		

Observations

- Justified using MP2 for fitting torsional parameters.
- Approximate functionals are getting better with time.
- DFT+D improves results in most cases.
- If a density functional better than MP2 exists, it isn't in any public code.
- CCSD(T) takes approximately 1 hour per job on 64 nodes.

CCSD(T)/cc-pVTZ energies at MP2/cc-pVTZ geometries.



This is about 100 years worth of computing...

Exascale computational chemistry

State of HPC for Chemistry

Both classical and ab initio molecular dynamics have essentially reached algorithmic maturity. Most research is fighting Amdahl's law and related concepts (FFT does not scale), e.g. DEShaw has turned classical molecular dynamics into an engineering problem.

Quantum many-body methods are far from algorithmic maturity because they have been constrained to tiny systems so the N-body problem is hidden behind dense linear algebra.

Dense linear algebra is great for Gordon Bell Prizes but terrible for science.

Deja Vu I (sort of)

February 15-19, 1999



Doubletree Hotel

2nd Conference on Enabling Technologies for Peta(fl)ops Computing Call for Participation and Papers February 15 - 19, 1999 Doubletree Hotel Santa Barbara, California

Santa Barbara, CA

Conference Chair: Paul Messina Caltech Program Chair: Thomas Sterling, Caltech/JPL Steering Committee Chair: Paul H. Smith, DOE

Sponsors: DARPA, NASA, NSF, DOE, NSA

The 2nd Conference on Enabling Technologies for Peta(fl)ops Computing is the first major open forum to treat the diversity of technical iss of in-depth workshops and sponsored studies conducted to explore the factors that will determine the ultimate path to realizing such capab understanding of Petaflops scale computing approaches and determine directions for future research leading to practical Petaflops perform of a wide range of issues and foster detailed discussion across conventional discipline boundaries. The conference will engage the interests areas associated with petaflops scale computing and beyond include but are not limited to:

Environmental Molecular Sciences Laboratory

Is this exciting? Yes!

- s Solution of longstanding problems
 - Detailed chemical reaction dynamics in solution
 - □ Molecular level environmental chemistry
 - □ Ab initio design of catalysts
 - Ab initio molecular dynamics
- s New "Grand Challenges"
- s Computational chemistry as a design tool

descriptions are a supplied to posed questions

Will the real impact of petaflop computing be that you can have all the tflop computing you need without waiting?On your desk?

Robert J. Harrison, 2/16/99

High Performance Computational Chemistry



Did we make good on our promises last time?

- Ab initio design of catalysts?
- Ab initio molecular dynamics?
- Molecular-level environmental chemistry?
- Detailed chemical dynamics in solution?

ARTICLES

Kemp elimination catalysts by computational enzyme design

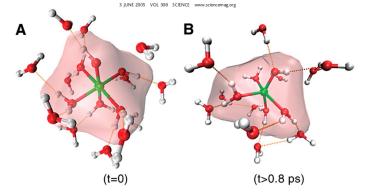
Daniela Röthlisberger!*, Olga Khersonsky¹*, Andrew M. Wollacott¹*, Lin Jiang^{1,2}, Jason DeChancie⁶, Jamie Betker³, Jasonie L. Gallaher³, Eric A. Althoff¹, Alexandre Zanghellini^{1,2}, Orly Dym³, Shira Albeck³, Kendall N. Houk⁶, Dan S. Tawfik¹ & David Baker^{1,2,3}

The design of new enzymes for reactions not catalysed by naturally occurring biocatalysts is a challenge for protein engineering and is a critical test of our understanding of enzyme catalysis. Here we describe the computational design of eight enzymes that use two different catalytic motifs to catalyse the Kemp elimination—a model reaction for proton transfer from carbon—with measured rate enhancements of up to 10° and multiple turnovers. Mutational analysis confirms that catalysis depends on the computationally designed active sites, and a high-resolution crystal structure suggests that the designs have close to atomic accuracy. Application of in vitro evolution to enhance the computational designs produced a >200-fold increase in $k_{\rm cat}/K_{\rm m}$ ($k_{\rm cat}/K_{\rm m}$ ($k_{\rm cat}/K_{\rm m}$), $k_{\rm cat}/k_{\rm mea}$ of $>10^{\circ}$). These results demonstrate the power of combining computational protein design with directed evolution for creating new enzymes, and we anticipate the creation of a wide range of useful new catalysts in the future.

Kinetic Evidence for Five-Coordination in AlOH(aq)²⁺ Ion

Thomas W. Swaddle, ¹ Jörgen Rosenqvist, ² Ping Yu, ³ Eric Bylaska, ⁶ Brian L. Phillips, ⁷ William H. Casey^{2,4,5}*

Trivalent aluminum ions are important in natural bodies of water, but the structure of their coordination shell is a complex unsolved problem. In strong and (pH = 3.0), AP exists almost writely at the octahedral All(F,0),e²⁺ ion, and (pH = 3.0), AP exists almost writely at the octahedral All(F,0),e²⁺ ion, in the shockmical layer of the strong and an artiflet of the structures are less clear. Other hydrolytic species, such as AlDH(a)[1-]; exist and are traditionally assumed to be heacondrate. We show, however, that the kinetics of proton and water exchange on aqueous AP, coupled with Car-Parmiello simulations, support a five-ocordinate. Mel show, however, that the kinetics of proton and water exchange on aqueous AP, coupled with Car-Parmiello simulations, support a five-ocordinate All(F,0)(AP) in an as the preclaminant from of AlChiquight under ambient conditions. This result exceeds the contraction of the co



Hard scaling challenges for *ab initio* molecular dynamics capabilities in NWChem: Using 100,000 CPUs per second

Eric J. Bylaska¹, Kevin Glass¹, Doug Baxter¹, Scott B. Baden² and John H. Weare³

¹ Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory, P.O Box 999, Richland, WA 99354

² Department of Computer Science and Engineering, University of California, San Diego

9500 Gilman Drive, # 0404, La Jolla, CA 92093-0404

³ Department of Chemistry and Biochemistry, University of California, San Diego 9500 Gilman Drive, # 0303, La Jolla, CA 92093-0303

Email: Eric.Bylaska@pnl.gov

Abstract. An overview of the parallel algorithms for ab intio molecular dynamics (AIMD) used in the NWChem program package is presented, including recent developments for computing exact exchange. These algorithms make use of a two-dimensional processor geometry proposed by Gygi et al. for use in AIMD algorithms. Using this strategy, a highly scalable algorithm for exact exchange has been developed and incorporated into AIMD. This new algorithm for exact exchange employs an incomplete butterfly to overcome the bottleneck associated with exact exchange term, and it makes judicious use of data replication. Initial testing has shown that this algorithm can scale to over 20,000 CPUs even for a modest size simulation.

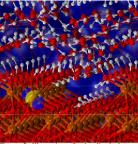


Figure 2. Illustration of a localized electron (i.e. polaron) on the surface of hematite calculated with a higher and more expensive level (e.g. hybrid DFT) of *ab initio* molecular dynamics. Lower levels of *ab initio* molecular dynamics will predict a delocalized electron.

Why you need to care about exascale.

The extrapolation fallacy

- None of you were computing before 1980.
- Frequency scaling provided continuous performance improvements from 1980-2005.
- Parallelism has replaced frequency-scaling as the source of improved performance.
- If you're code isn't fully SIMD and multithreaded, it is going to run slower with time.

Possible solutions

- Pay a lot of money for processors aimed at Wall Street (high-frequency traders). Stock up now.
- Stop using legacy codes that aren't properly parallel. Switch to modern parallel codes and become a supportive memory of their user community.
- Learn to write massively-parallel code. Produce your own open-source package and pray to the Flying Spagetty Monster that people use it and you get continuous funding.

Future trends

There are three basic architectural trends you need to be aware of:

- cloud computing: it runs "out there somewhere" (probably not so quickly).
- GPU computing: requires near-complete rewrite of application code to improve more than 2-3 times.
- supercomputing: both GPU-based and 1M-core systems are just around the corner.

Next time...

- Ensemble parallelism on Blue Gene/P.
- GPU quantum chemistry.
- Overview of free supercomputer resources and how to get access to them.
- Please suggest other topics of interest.

Acknowledgments





Argonne: Nick Romero, Pavan Balaji, Jim Dinan, Vitali Morozov,

Eugene DePrince

PNNL: Karol Kowalski, Niri Govind

UChicago: Benoît Roux, Karl Freed, Glen Hocky

ORNL: Vinod Tipparaju

Summer Students: Jack Poulson (UTexas), Sreeram Potluri

(OSU), Kevin Stock (OSU)

